UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

God is in the rain: The impact of rainfall-induced early social distancing on COVID-19 outbreaks.

Permalink

<https://escholarship.org/uc/item/3jr4k0cm>

Authors

Shenoy, Ajay
Sharma, Bhavyaa
Xu, Guanghong
et al.

Publication Date

2022

DOI

10.1016/j.jhealeco.2021.102575

Peer reviewed

Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.

God is in the rain: The impact of rainfall-induced early social distancing on COVID-19 outbreaks

Ajay Shenoy^{a,1,*}, Bhavyaa Sharma^a, Guanghong Xu^a, Rolly Kapoor^a, Haedong Aiden Rho^a, Kinpritma Sangha^b

^a University of California, Santa Cruz, United States

^b Siemens Healthineers

ARTICLE INFO

JEL classification:

I12
I18
H75

Keywords:

COVID-19
Coronavirus
Social distancing
Rainfall

ABSTRACT

We measure the benefit to society created by preventing COVID-19 deaths through a marginal increase in early social distancing. We exploit county-level rainfall on the last weekend before statewide lockdown in the early phase of the pandemic. After controlling for historical rainfall, temperature, and state fixed-effects, current rainfall is a plausibly exogenous instrument for social distancing. A one percent decrease in the population leaving home on the weekend before lockdown creates an average of 132 dollars of benefit per county resident within 2 weeks. The impacts of earlier distancing compound over time and mainly arise from lowering the risk of a major outbreak, yielding large but unevenly distributed social benefit.

1. Introduction

When COVID-19 first reached U.S. shores, local governments differed in how quickly they imposed mandatory restrictions on businesses and gatherings. California began its statewide lockdown 3 days earlier than New York. Even within California, seven counties in the San Francisco Bay Area began their lockdown a few days earlier than the rest of the state. These early-movers had smaller outbreaks in the initial phase of the pandemic, leading many to suggest earlier lockdowns were a key driver of their success (Amuedo-Dorantes et al., 2020; Dave et al., 2020b,a). But lockdowns and other less extreme measures have generated enormous controversy because of their costs to the economy (Baek et al., 2020; Cheng et al., 2020; Coibion et al., 2020; Fairlie, 2020; Gupta et al., 2020a; 2020b) and to physical and mental health (Altindag et al., 2020; Leslie and Wilson, 2020; Ziedan et al., 2020). Judging whether a faster response was justified requires weighing these costs against the social benefit of the lives saved. Though the costs are visible, the benefit of a marginal increase in earlier social distancing has not been rigorously measured.

Any such measurement requires careful estimates of the number of deaths averted by earlier social distancing. A naïve comparison between states risks conflating the impact of earlier distancing with differences in state characteristics. Fig. 1 shows that states that issued earlier lockdowns have higher median incomes and more college degree holders, but fewer black and older residents—factors shown to be correlated with transmission and death rates (Allcott et al., 2020; Knittel and Ozaltun, 2020). Even within a state, locales

* Corresponding author at: Rm. E2-455, University of California, M/S Economics Department, 1156 High Street, Santa Cruz CA, 95064.

E-mail address: azshenoy@ucsc.edu (A. Shenoy).

URL: <http://people.ucsc.edu/~azshenoy> (A. Shenoy), <https://economics.ucsc.edu/about/people/phd.html> (B. Sharma), <https://economics.ucsc.edu/about/people/phd.html> (G. Xu), <https://economics.ucsc.edu/about/people/phd.html> (R. Kapoor), <https://economics.ucsc.edu/about/people/phd.html> (H.A. Rho), <https://www.linkedin.com/in/kinpritma-sangha-31100b92/> (K. Sangha)

¹ We are grateful to the editor Dr. Kosali Simon and three anonymous referees for thoughtful suggestions. We thank SafeGraph for providing access to their data. We also appreciate helpful comments from Laura Giuliano, Justin Wolfers, and innumerable individuals on Twitter.

Fig. 1. States that lock down earlier are systematically different on baseline characteristics.

Note: The size of each circle is proportional to the number of states that shut down on that date. Demographics are from the 2014–2018 American Community Survey (5-year estimates). Dates of state-wide lockdown orders come from the Institute of Health Metrics and Evaluation. See [Section 2.1](#) for details about the data.

that issued earlier lockdowns may differ systematically in ways that may or may not be observable. For example, the Associated Press reports that the Bay Area lockdown had its roots in an association of local health officials that formed during the AIDS epidemic and has met regularly to discuss prior epidemics like Ebola and swine flu ([Rodriguez, 2020](#)). The presence of such an institution may have had other impacts on the local response to COVID-19 beyond the lockdown. It is also possible that invisible factors like social trust ([Bartscher et al., 2021](#)), civic capital ([Barrios et al., 2021](#)), culture ([Bazzi et al., 2021](#)), and trust in institutions ([Bargain and Aminjonov, 2020](#)) drive both private responses to the pandemic and political appetite for lockdowns, making it difficult to isolate the effect of early social distancing. The problem of selection bias is compounded by the problem of measurement. It is possible that the states and counties that responded more quickly are also more active in testing for the disease, creating non-classical measurement error. Given these confounders, it is not surprising that simple within-state comparisons produce nonsensical results.²

We sidestep these challenges by exploiting within-state variation in early social distancing induced by rainfall. We measure county-level rainfall on the last weekend before the county's home state went into mandatory lockdown. This key weekend is the last time that people had wide discretion in leaving home for reasons unrelated to work (dining at restaurants, for example). Focusing on this weekend creates a natural experiment for a marginally longer period of social distancing. After controlling for average historical rainfall, temperature, and state fixed effects, rainfall on this specific weekend is plausibly exogenous. Counties that had heavy rainfall were exogenously induced to exercise a marginal degree of extra social distancing just a few days before counties that had less rainfall. We measure how many fewer COVID-19 cases and deaths these rainy counties had in the weeks after the statewide lockdown. Finally, we put these estimates into a metric that is directly comparable to the economic cost of a lockdown. We combine our estimates of deaths averted with several recent estimates of the value of a statistical life to calculate the average per capita benefit of a marginal change in the degree of early social distancing.

We estimate that marginal additional distancing bends the trajectory of a local outbreak, with benefits still accumulating even two weeks after the statewide lockdown, many days after the crucial weekend. The two-stage least squares estimates imply that a 1 percentage point increase in the number of people leaving home causes an additional 14 cases and 1.4 deaths per 100,000 residents. These effects are all the more remarkable because the variation in social distancing induced by rainfall, though precise, is relatively small. But the impact of the initial reduction is propagated over time. We measure growing impacts that have not leveled off even

² We show in Appendix A.2 that ordinary least squares estimates actually have the wrong sign, implying greater social distancing predicts more infections.

18 days after the lockdown, nearly 3 weeks after the crucial weekend. These effects appear to be driven by the right tail of the distribution. Counties where more people left home on the pre-shutdown weekend are no more likely to have a marginally higher case count, but are slightly more likely to have a big outbreak. This result is what might be expected given that differences in the number of infections on the eve of a statewide lockdown will either vanish or be drastically amplified depending on whether the county lowers the viral reproduction rate below 1 and avoids “superspreader” events.

We calculate the dollar value of deaths averted by a marginal increase in early distancing by leveraging several estimates in the literature of the value of a statistical life. A commonly cited value from the U.S. Environmental Protection Agency implies that a 1 percent reduction in people leaving home on the weekend before the state-wide lockdown yields a per-capita benefit of 132 dollars for the average county resident within 14 days of the statewide lockdown. A more conservative estimate that adjusts for the age profile of COVID-19 deaths halves the per capita value, though even this smaller amount exceeds the earnings from 8 hours of work at the federal minimum wage. Measuring benefits at the longer horizon of 18 days raises the value by roughly one-third, suggesting that a marginal increase in early distancing is an investment that pays sizeable returns. And since we account only for the value of deaths averted, the full value—which should include reduced healthcare costs and reduced incidence of chronic illness—is undoubtedly higher than our estimates imply. This exercise reveals that, although putting dollar values on lives may seem callous, it is only by doing so that the sheer magnitude of the benefit becomes apparent.

Our paper follows many studies in economics that exploit natural experiments to measure the impact of public health interventions on mortality. Prior work has studied the impact of community health centers (Bailey and Goodman-Bacon, 2015); campaigns to vaccinate against influenza (Ward, 2014) and reduce alcohol consumption (Bhattacharya et al., 2013); access to public health insurance during a flu pandemic (Clay et al., 2020); packaging and messaging of Malaria drugs (Cohen and Saran, 2018), and the discovery of antibiotics (Jayachandran et al., 2010). Other work has studied the impact of disease eradication on long-run health and economic outcomes (e.g. Acemoglu and Johnson, 2006; Baird et al., 2016; Bleakley, 2003; 2010; Hamory et al., 2020). Our work is among the first to study social distancing, an intervention that requires no special technology but whose economic costs have been hotly contested.

Our paper also joins a small but growing number of studies that provide guidance for weighing the costs and benefits of policies that can save lives in the COVID-19 pandemic. Most of these studies infer the number of deaths averted by using epidemiological models (Acemoglu et al., 2020; Greenstone and Nigam, 2020; Thunstrom et al., 2020). Our approach is novel in exploiting a natural experiment to estimate deaths averted by marginal changes in the degree of early social distancing. In making these estimates our study also joins a growing literature on the impact of social distancing on COVID-19 transmission. Pei et al. (2020) use an epidemiological model to simulate COVID-19 trajectories in a counterfactual world where lockdowns had begun a few weeks sooner. Our study approaches this question using a natural experiment rather than a model. A few recent studies (Courtemanche et al., 2020; Fowler et al., 2020; Spiegel and Tookes, 2020) use difference-in-differences designs to study the impact of statewide closures and lockdowns on transmission. Aside from exploiting an orthogonal source of variation, our study aims to answer a different question: whether marginal improvements in early distancing can affect medium-run outcomes. Meanwhile, Brzezinski et al. (2020) use state-level rainfall and temperature as exogenous variation in non-mandated social distancing to study whether state governments are less likely to mandate social distancing where it is already being practiced.³ Finally, our study joins a small number of studies that uses instrumental variables to test whether COVID-19 outbreaks are mitigated by public health measures, notably masks (Welsch, 2020) and remote work (Glaeser et al., 2020; McLaren and Wang, 2020).

Methodologically our study is most similar to Madestam et al. (2013), which measures the impact of rainfall on a single pivotal date (Tax Day 2010) to measure the long-run impacts of Tea Party protests. One major advantage to studying a one-time shock rather than panel variation is that we can fully trace differential trajectories across counties. And since that shock is on the weekend before statewide lockdown, it is the closest possible counterfactual to having a longer policy of social distancing.

Our results suggest that even marginal additional distancing in the early stages of an outbreak have large economic benefit. These estimates contribute to the retrospective discussion of whether aggressive policies in the early stages of the pandemic were justified, and whether similar measures would be justified in the face of either a resurgence of the coronavirus or a future respiratory illness. Social distancing measures have generated enormous controversy because of their highly visible costs. The broader public discourse has at times grown distant from the evidence, making a sober analysis of the costs and benefits all the more crucial.

2. Research design

The key input into our calculation is an unbiased estimate of the impact of marginal earlier distancing on COVID-19 cases and deaths. Our ideal experiment would be to randomly assign some counties to begin social distancing sooner than others. Since such an experiment is not feasible, our natural experiment focuses on rainfall-induced social distancing on the weekend just prior to the statewide lockdown. People in rainy counties began a marginal degree of additional social distancing a few days sooner than other counties.⁴

We apply the method of instrumental variables. We use rainfall on the last weekend before statewide lockdown as an instrument for the percentage of people leaving home that weekend, as measured from mobile phone data. We control for historical rainfall (average rainfall on this calendar weekend over the prior 5 years) as well current temperature and historical temperature.

³ Since we exploit only within-state variation, their result is not a threat to our design. We verify in Section 3.4 and Appendix A.8 that county-level policy responses do not bias our results.

⁴ The gap between the last weekend (as defined in Appendix 2) and the shutdown is 3 days for the median county.

Fig. 2. First-stage and reduced form.

Note: Each panel shows a partial correlation plot of rainfall on the weekend before the statewide lockdown against either the percentage of people leaving home on that weekend (left-hand panel) or total cases per 100,000 as of 14 days after the lockdown. We calculate residuals from a regression of both X and Y variable on state fixed-effects, historical rainfall, current and historical temperature, and baseline case controls. We define bins based on residualized rainfall. Each dot shows the average residualized outcome within the bin, and the line shows the linear prediction. The histogram shows the number of observations that fall into each bin.

2.1. Data

Weather: We measure average precipitation, which we will colloquially refer to as “rainfall,” by spatially merging weather stations from the Global Historical Climatology Network-Daily Database (Menne et al., 2012) to U.S. counties based on 2012 Census TIGER/Line shapefiles. We calculate county-level average precipitation and daily maximum temperatures. For each day in 2020 we calculate the average precipitation and max temperature for that same day-of-year from 2015–2019. We then take the inverse hyperbolic sine of all of these quantities.⁵ This transformation is a standard way to approximate a log transformation without having to discard zero-precipitation days. As long as precipitation itself is exogenous, the transformed quantity is also exogenous and, as we show in Fig. 2 below, has a roughly linear relationship with our primary measure of social distancing. We apply the same transformation to temperature to maintain consistency. From here on we refer to these transformed quantities as simply current or historical rainfall and temperature.⁶

Social Distancing: Our primary measure of social distancing is the percentage of people that leave home, calculated using aggregated mobile phone GPS data provided by SafeGraph (SafeGraph, 2020a).⁷ The data report the total devices in SafeGraph’s sample by block group, and the number that leave their home.⁸ We aggregate these two counts by county and calculate the percentage leaving home.

Leaving home is our first-stage regressand because keeping people home is the primary impact of rain on social distancing, and keeping people at home for an extra weekend is the most natural analogy to locking down a few days sooner. But to better understand what activities people are deterred from doing when they stay home—and whether those who do leave change where they go—we draw on several other measures of social distancing. We use two measures of indoor exposure. The first is the Device Exposure Index (Couture et al., 2020a), which represents the number of people (cell phones) an average individual was exposed to in small commercial venues within the county. We also use SafeGraph’s Weekly Patterns data to compute a measure of “gatherings” based

⁵ The inverse hyperbolic sine transformation $\log(x + \sqrt{x^2 + 1})$ is a convenient approximation to the natural logarithm that is well-defined when $x = 0$ and converges to $\log 2 + \log x$ as $x \rightarrow \infty$.

⁶ Though the inverse hyperbolic sine transformation is not invariant to units of measurement (unlike the log), that limitation is less problematic for us because the transformed variable is only used as an instrument. The estimator is still valid as long as the choice of units does not cause the instrument to become endogenous, which is highly unlikely.

⁷ SafeGraph is a data company that aggregates anonymized location data from numerous applications in order to provide insights about physical places. To enhance privacy, SafeGraph excludes census block group information if fewer than five devices visited an establishment in a month from a given census block group.

⁸ SafeGraph defines “home” as the “common nighttime location of each mobile device over a 6 week period to a Geohash-7 granularity (153m x 153m).” Leaving home is defined as leaving that square. Though this measure of staying home is imperfect, two-stage least squares will prevent measurement error from biasing the estimates.

on whether more than 5 devices ping within a single indoor non-residential location within one hour (SafeGraph, 2020b). Since the SafeGraph sample represents roughly 6% of a typical county, 5 devices represent a large number of people. We rescale both measures by their daily average on the first full weekend in March, meaning a value of 100 denotes the same exposure or number of gatherings as the first weekend of March (which was before any local or state lockdown).

We also use several measures of long-distance travel. Using SafeGraph's data we measure the percentage of devices that travel greater or less than 16 km from home (among those that leave home). We also measure cross-county travel using the Location Exposure Index (Couture et al., 2020b). We measure the fraction of people in a county who were not present on any of the prior 14 days.⁹

COVID-19 Cases and Deaths : We measure daily (cumulative) COVID-19 cases/deaths by combining data from Johns Hopkins University and the CoronaDataScraper project (Center for Systems Science and Engineering (Johns Hopkins University) (2020); Corona Data Scraper (2020)). As described in detail in Appendix B, we manually corrected missing values by consulting county public health departments and local newspapers. All of these measures are cumulative cases and deaths rather than new cases and deaths. Our primary outcomes are the number of cases and deaths per 100,000 population, measured 14 days after the statewide lockdown (which we will refer to as “endline”).¹⁰

Demographics : We measure demographic characteristics (such as population size, median income, age profiles of the population) using the 2014–2018 five-year estimates from the American Community Survey (Manson et al., 2019).

Lockdowns: Finally, we measure statewide lockdown dates using the Institute of Health Metrics and Evaluation's record of state policies as of 17 April 2020 (Institute for Health Metrics and Evaluation, 2020). The dataset has all shutdown dates up to 7 April. Any state that had not shut down by that date (or was not recorded as doing so by the Institute) is excluded from our study.

2.2. Instrument and specifications

Defining the Instrument: We identify the last Saturday and Sunday before the day of the lockdown order. If the lockdown began on a Sunday we take only the Saturday of that weekend as the “weekend before.” If it begins on a Saturday we take the prior weekend. We average the transformed values of rainfall and temperature (both current and historical) as well as social distancing across the days of this weekend. We compute baseline cases and deaths as those recorded for the day before this last weekend, and baseline growth in these measures as the average change in the inverse hyperbolic sine of each in the prior 7 days.

Specification, Identification, and Inference: We estimate first-stage, reduced form, and second-stage regressions of the form

$$D_i = \alpha_s + \gamma R_i + \tau_1 \bar{R}_i + \tau_2 T_i + \tau_3 \bar{T}_i + X_i \omega + u_i \quad (1)$$

$$Y_i = \zeta_s + \rho R_i + \xi_1 \bar{R}_i + \xi_2 T_i + \xi_3 \bar{T}_i + X_i \theta + v_i \quad (2)$$

$$Y_i = \kappa_s + \beta \hat{D}_i + \phi_1 \bar{R}_i + \phi_2 T_i + \phi_3 \bar{T}_i + X_i \vartheta + z_i \quad (3)$$

where i and s index counties and states, D is the percentage of people leaving home, Y is the outcome, α_s and κ_s are state fixed-effects, R and $\bar{R}$ are current and historical rainfall, T and $\bar{T}$ are current and historical temperature, and X is a vector of baseline and demographic control variables that vary across specifications, with the most basic specification having no controls. Since there is spatial correlation in both rainfall and COVID-19 infections, we cluster standard errors using a $3^\circ \times 3^\circ$ latitude-longitude grid.¹¹

We must control for historical rainfall because even within a state, counties that are typically rainy in March and April may be systematically different from those that are not (e.g. Santa Cruz, California versus San Diego). The instrument R_i is thus excess or unexpected rainfall, which is plausibly uncorrelated with historical demographic characteristics. We control for temperature because some experts and politicians have hypothesized that it may directly impact COVID-19 transmission.¹²

The identification assumption is that, after controlling for state fixed-effects, historical rainfall, and temperature, rainfall on the pre-shutdown weekend only affects endline case counts through its impact on the number of people leaving home. Under this assumption, β is the average treatment effect on the outcome of a 1 percentage point increase in the number of people leaving home on the weekend before the lockdown.¹³

In reality, the assumption of the classical IV model—that rainfall uniformly reduces everyone's probability of leaving home, regardless of which activities they were planning—may not hold. It is unlikely that anyone is *more* likely to leave home when it

⁹ For more information on the Device Exposure Index and the Location Exposure Index see Appendix B.1.

¹⁰ We choose these measures both because they are the measures most commonly used by policymakers to gauge the severity of an outbreak, and because they give the most accurate reflection of the number of infections relative to the number who could potentially be infected. We choose 14 days as our default horizon because this is the typical quarantine period for the disease, though Section 3.2 shows the impact at every horizon.

¹¹ To be precise, we generate a grid where each cell is of length 3° on each side. Each county is uniquely assigned to the cell that contains its centroid. Clustering within grid cell allows for arbitrary correlation in the regression error term of counties within the same cell. The median cell contains 8 counties, though the smallest contains 1 and the largest contains 60. The number of counties will vary based on their size and remoteness (counties are generally smaller in the eastern U.S., for example). We chose this method of clustering rather than, say, state-level clustering because many states, especially in New England, are geographically tiny and neither rainfall nor a coronavirus outbreak will remain contained within state boundaries. Though the precise size of the cluster is somewhat arbitrary, we show in Appendix A.10 that the statistical significance remains intact even with far smaller and far larger clusters.

¹² Chin et al. (2020), for example, find that temperature affects virus stability in lab samples.

¹³ Since there is a single endogenous regressor and a single excluded instrument, $\hat{\beta} = \hat{\rho}/\hat{\gamma}$.

is raining—there are no defiers, to use the language of local average treatment effects (Angrist and Imbens, 1995; Angrist et al., 1996). The compliers—those who stay home because of rainfall—will be some combination of people who were planning purely outdoor activities, people who were planning indoor activities outside the home but prefer not to travel in the rain, and people who were planning indoor activities bundled with outdoor activities (e.g. a day at the beach and dinner at a restaurant). Since the group deterred from purely outdoor activities is most likely to be moved by the instrument despite never being at high risk of infection (Bulfone et al., 2020), we would expect the local average treatment effect to be an *underestimate* compared to the average treatment effect.

It is also possible that some of the people who choose to leave home despite the rain (the never-takers) actually substitute from outdoor to indoor activities (dining indoors instead of outdoors, for example). To the extent that this happens it would further imply that we are actually underestimating the impact of early distancing on COVID-19 transmission. As we show in Section 3.1, the net impact of rainfall on our measure of indoor exposure is negative, suggesting the reduced exposure of those who stay home because of the rain outweighs any substitution towards indoor activities by those who leave home despite it.

Additional Control Variables: Since rainfall is exogenous, the control variables X_i will not affect the consistency of the estimates. But they can make the estimates more precise by reducing the unexplained variation in social distancing and COVID-19 cases and deaths. Our basic specification includes nothing in X_i . Our preferred specification adds controls for baseline COVID-19 prevalence. We include the number of cases per 100,000 at baseline, the raw number of cases at baseline, and the growth rate of cases in the week prior to the pre-lockdown weekend.¹⁴ Our most comprehensive specification includes baseline controls as well as demographic characteristics.¹⁵

3. Results: Impact of earlier distancing on COVID-19 cases and deaths

3.1. Basic estimates

First-Stage—Impact of Rainfall on Social Distancing: Column 1 in Panel A of Table 1 shows estimates of the first-stage (Eq. 1). The coefficient of 0.4 implies that after controlling for historical rainfall and temperature, a 10 percent increase rainfall from its sample median causes a 0.036 percentage point decrease in the number of people who leave home.¹⁶ The estimate implies that moving a county from zero rainfall to the 90th percentile of the distribution would cause a 1.8 percentage point reduction in the number of people leaving home—comparable to the reduction caused by a stay-at-home order (see Section 5).¹⁷ The F-statistic is 11.82, meeting typical standards for instrument strength, but in our two-stage least squares results below we also report weak identification-robust tests for significance.

Although some people will stay home because of the rain, those who do leave may be more likely to pack into bars and restaurants instead of visiting the outdoors. Although this would imply the IV coefficient would underestimate the impact of staying at home, we can test directly whether switching to indoor activities by never-takers outweighs the reduced exposure of the compliers. Column 2 shows that rainfall causes average indoor exposure, based on how many people visit small indoor venues, to decline by 0.9 percentage points relative to its level the first weekend of March (prior to any lockdown). Column 3 shows that our measure of large gatherings declines by 1.7 percentage points relative to early March. Taken together, these estimates suggest that any potential substitution towards indoor activities is not large enough to outweigh the first-order impact of staying at home.

One potential explanation is that rainfall makes people less likely to travel. If the main reason for travel is an outdoor activity (visiting the beach) that also requires indoor activities (eating at restaurants and staying at a hotel), it may in part explain why net indoor exposure falls. Columns 4 and 5 measure the impact on the percentage of people leaving home and traveling a short or long distance (based on whether they traveled more than 16 km from home). The estimates suggest a larger impact on long distance travel (especially compared to the mean). Column 6 shows that a one-unit increase in rainfall causes a 0.28 percentage point decrease in the fraction of people in the county who had not been there in the previous two weeks, suggesting a sizable decline in cross-county travel.

Reduced-Form and Two-Stage Least Squares: Panel B of Table 1 shows estimates of the reduced-form impact of rainfall on COVID-19 cases and deaths per 100,000 at endline, which these regressions define as 14 days after the statewide lockdown. Column 1 shows that a 1 unit increase in rainfall on the weekend before lockdown lowers the number of cases at endline by 6.8 per 100,000. Columns 2 and 3 show that controlling for baseline prevalence and demographics tightens the standard errors without substantially changing the estimates. Columns 4—6 imply that the reduction in cases translates to a reduction in deaths, as well. A 1 unit increase in rainfall causes a 0.5 to 0.7 per 100,000 reduction in the death rate.

Fig. 2 shows a partial correlation plot of the first-stage and reduced form of the regression in Column 2 (which includes baseline case controls). The plot illustrates how rainfall on the last weekend before the state-wide lockdown lowers both the percentage of

¹⁴ We control for both cases per 100,000 and raw case counts at baseline because both are independently informative about social distancing and endline outcomes. That is likely because while the one measures the baseline rate of prevalence, the other drives initial local media coverage. It is also likely that a greater raw number of cases lowers the probability that the infection dies out (say, if all initially infected people self-isolate). The case growth rate, which we calculate as the average change in the inverse hyperbolic sine of case counts, is informative about the trajectory prior to the pre-shutdown weekend.

¹⁵ Total population; fraction of population in the bins 60–69, 70–79, and over 80; fraction African American; and median household income.

¹⁶ The median is 1.85 tenths of a millimeter, and $\sinh^{-1}(1.1 \times 1.85) - \sinh^{-1}(1.85) \approx 0.085$.

¹⁷ The 10th percentile is 0, and the 90th percentile is 35 tenths of a millimeter.

Table 1
Two-stage least squares estimates.

Panel A: Interpreting the First-Stage						
	First-Stage	Activities Averted by Staying Home				
	(1)	(2)	(3)	(4)	(5)	(6)
	% Leaving Home	Exposure	Gatherings	Travel Near	Travel Far	Non-Locals
Rainfall	−0.434*** (0.126)	−0.905*** (0.330)	−1.710** (0.691)	−0.216* (0.130)	−0.340** (0.145)	−0.278** (0.122)
Counties	1951	1399	1762	1951	1951	1399
Clusters	139	113	124	139	139	113
Outcome Mean	64.78	37.38	38.35	41.41	21.43	9.14
F-stat: Rainfall	11.82	7.54	6.13	2.75	5.51	5.17
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Panel B: Reduced-Form						
	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	−6.751** (3.146)	−6.106*** (1.700)	−5.961*** (1.657)	−0.719 (0.461)	−0.583*** (0.214)	−0.557*** (0.204)
Counties	1951	1951	1951	1951	1951	1951
Clusters	139	139	139	139	139	139
Outcome Mean	58.27	58.27	58.27	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X
Panel C: Two-Stage Least Squares						
	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	15.566 (9.554)	14.467*** (4.806)	16.737*** (5.845)	1.658 (1.266)	1.381** (0.552)	1.564** (0.638)
Counties	1951	1951	1951	1951	1951	1951
Clusters	139	139	139	139	139	139
First-Stage F	11.82	16.78	19.84	11.82	16.78	19.84
Anderson-Rubin Chi2	0.03	0.00	0.00	0.11	0.01	0.01
Stock-Wright S-Stat	0.00	0.00	0.00	0.00	0.00	0.00
Outcome Mean	58.27	58.27	58.27	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: All standard errors are clustered using a 3°x 3° latitude-longitude grid to adjust for spatial correlation. **Panel A:** “Exposure” refers to the Device Exposure Index, a measure of the number of devices (cell phones) visiting small indoor venues. “Gatherings” measures the number of times more than 5 devices ping in a single indoor venue within the span of an hour. Both of these measures are rescaled as a percentage of their level on the weekend 7–8 March. “Travel Near” and “Travel Far” give the percentage of devices that leave home and travel less than versus more than 16 km. “Non-Locals” gives the percentage of devices in the county that were not present on any of the prior 14 days. **Panels B and C:** “Baseline Case Controls” are the number of COVID-19 cases the day before the pre-shutdown weekend (both the raw count and the number per 100,000), and the average growth (change in the inverse hyperbolic sine) of cases in the week preceding the last weekend. “Demographic Controls” are total population; fraction of population in the bins 60–69, 70–79, and over 80; fraction African American; and median household income. “Anderson-Rubin Chi2” and “Stock-Wright S-Stat” give the p-values from the weak identification-robust tests of [Anderson and Rubin \(1949\)](#) and [Stock and Wright \(2000\)](#) for the statistical significance of the IV estimate. *p=0.10 **p=0.05 ***p=0.01

Fig. 3. The excess case count in counties with less early distancing continues to increase even 18 days after lockdown.

Note: Using total cases per 100,000 at each horizon $h = 2, 4, 6, \dots, 18$ we estimate the two-stage least squares coefficient controlling for baseline case controls (analogous to Column 2 of Panel C, Table 1). Each coefficient is from a separate regression (and the regression at $h = 14$ is identical to that reported in Table 1).

people leaving home (left-hand panel) and the number of cases at endline (right-hand panel). The plot shows that our estimates are not driven by outliers, and that both relationships are approximately linear. The linearity is a crucial check to confirm that the way we measure rainfall produces a valid first-stage.

Under the assumption that rainfall only affects disease transmission through its impact on early social distancing, the two-stage least squares estimate—the ratio of the reduced-form and first-stage coefficients—gives the causal impact of early social distancing on COVID-19 cases and deaths. Panel C of Table 1 presents these estimates. All three specifications have a strong first-stage, with the F-statistic on the excluded instrument (weekend rainfall) varying from 11 to 18. The basic specification, which has no controls, is relatively noisy and statistically insignificant.

But after controlling for baseline case controls the standard errors become tight enough to make the estimates highly significant (Columns 2 and 5). The final specification additionally controls for county demographics, which makes little difference in size or significance of the estimates (Columns 3 and 6). Indeed, all three specifications produce near-identical estimates. A 1 percentage point increase in the number of people leaving home on the weekend before the shutdown causes an additional 14 to 17 cases and 1.4 to 1.7 deaths per 100,000. The Anderson-Rubin and Stock-Wright tests, which have superior size properties compared to t-ratio-based inference, confirm that these estimates are highly significant.¹⁸

The size of these estimates relative to the mean of the outcome may seem surprising. But it is worth recalling that a large fraction of counties in our sample, drawn from the early days of the pandemic, had zero confirmed cases of COVID-19. It is more useful to benchmark the estimates against the population that is susceptible to infection, which is roughly the entire population during this period of the pandemic. By this benchmark our results imply that an additional 1 percentage point of people leaving home causes an additional 0.014 percentage points of the population to become infected and 0.001 percentage points to die. These estimates are not unreasonable given that COVID-19 is a highly infectious respiratory virus.

3.2. Comparative dynamics in counties with less early social distancing

Table 1 gives a relatively limited picture of the trajectory of cases because all outcomes are measured at the fixed horizon of 14 days after the statewide lockdown. One advantage of our research design is that we can estimate the comparative dynamics of case rates between counties that quasi-randomly practiced different levels of early social distancing. Using the same specification as Columns 2 and 5 of Table 1.C, we estimate the impact on cases and deaths per 100,000 2 days after the lockdown, 4 days after, and so on for every horizon $h = 2, 4, 6, \dots, 18$. Fig. 3 plots each coefficient against h . The estimated impact appears to increase linearly over time with no sign of leveling off within the horizon available to us.¹⁹ The figure suggests the impact of a one-time difference in early social distancing is surprisingly long-lived. Cases and deaths show a similar pattern, though deaths increase at a somewhat

¹⁸ Many studies (for a review, see Andrews et al., 2019; Lee et al., 2021) question the common rule of thumb that two-stage least squares t-statistics are valid when the first-stage F-statistic exceeds 10. The tests proposed by Anderson and Rubin (1949) and Stock and Wright, 2000 are properly-sized even in the presence of weak instruments.

¹⁹ At longer horizons we would start to lose states because our case count data ends 18 days after the last state in our sample to go on lockdown.

Fig. 4. Counties with less social distancing are more likely to have very large (right-tail) outbreaks.

Note: We estimate the impact across the distribution of outcomes. Each point and confidence interval is the two-stage least squares estimate of the impact of early social distancing on the probability of having endline cases per 100,000 greater than the percentile or absolute number indicated on the horizontal axis. Each estimate controls for baseline case rate, count, and growth (analogous to Column 2 of Panel C, Table 1).

slower pace. It may seem surprising that there are impacts on reported cases and deaths at a horizon of even 4 days, but recall that 4 days after the statewide lockdown is on average a full week after the weekend of transmission.

We find no evidence, however, that the growth *rate* of cases increases because of more people leaving home on the last weekend (see Appendix A.5). That is not surprising because the natural experiment induces some counties to begin early social distancing just before all counties go uniformly into lockdown. The effect is analogous to quasi-randomly inducing some counties to begin lockdown with a larger infected population. As long as this difference in initial population does not affect how carefully the lockdown is observed, it will rescale the case count without affecting the transmission rate.²⁰

3.3. Distributional impact: Early social distancing lowers the chance of right-Tail outcomes

Given the nature of exponential growth, local COVID-19 outbreaks may quickly die down or rapidly spiral out of control. Since middling outbreaks are unlikely, early social distancing might mainly affect the distribution of outcomes by reducing weight in the right tail rather than lowering the median. We test for the impact on the full distribution by defining dummies for whether the endline number of cases per 100,000 is greater than each decile of the distribution. We estimate Eq. (3) using these dummies as the outcomes (using the specification with baseline case controls). This procedure is analogous to testing how the inverse cumulative distribution function is shifted by a 1 percentage point reduction in early social distancing.

The left-hand panel of Fig. 4 plots the estimates with their 90 and 95 percent confidence intervals. The figure suggests that although the estimated impact becomes positive around 0.4 (meaning less early distancing increases the probability of being above the 40th percentile), the effect only becomes significant at 0.7. That suggests early distancing is lowering the probability of a right-tail outbreak. The most precise estimate is the last. A 1 percentage point increase in the number of people leaving home on the weekend before lockdown causes a 2 percentage point increase in the probability of an outbreak that puts the county in the top 10 percent of the distribution. The right-hand panel clarifies just how large these right-tail events are. This panel is analogous to the first one, but it defines dummies based on having a case rate at endline (14 days after statewide lockdown) above some absolute cutoff. The size and significance peaks at 100 cases per 100,000, a very large case count.

The results suggest early social distancing worked less by causing a moderate reduction in cases than by reducing the chance of a big outbreak. This result may be consistent with several recent studies that find that COVID-19 has a very low dispersion factor, meaning small groups of “superspreaders” are responsible for the vast majority of cases (Kupferschmidt, 2020). Endo et al. (2020) estimate using a mathematical model that as few as 10% of initially infected people may be responsible for as much as 80% of subsequent cases. Miller et al. (2020) find a similar result when they use genome sequencing to trace the virus’s spread across Israel. If early social distancing marginally reduces the probability a superspreader begins a transmission chain, it could explain why our estimates

²⁰ If endline case count is $Y_T = Y_0 \exp(gT)$, our natural experiment is analogous to increasing Y_0 .

are driven by changes in the number of large outbreaks. Regardless of the cause, our estimates imply that most counties that began distancing sooner had little benefit, but those that did benefit did so tremendously.

3.4. Summary of robustness checks

In the appendix we show the results of several other tests:

Balance: One concern is that rainfall, even after controlling for state fixed effects, historical rainfall, and current and historical temperature, is not truly exogenous. We show in Appendix A.1 that rainfall is uncorrelated with baseline measures of COVID-19 prevalence and county demographic characteristics.

Heterogeneity: We show in Appendix A.3 that there is little evidence of heterogeneous impacts by baseline case levels, baseline case growth, the time between the last weekend and the start of the statewide lockdown, and a host of county-level demographic characteristics. This seems largely a consequence of not having enough data to generate a strong first stage when splitting the sample or identifying an interaction as well as a direct effect. There is some slight evidence that early social distancing has less of an impact in counties with an older population, though the mechanism for that result is uncertain.

Outliers: Given that Section 3.3 shows the effect comes largely from changes in the likelihood of right-tail events, one may worry that the entire estimate is driven by a few outliers. Appendix A.4 shows that winsorizing the very largest outcomes still yields significant effects. Although the top of the distribution does drive the result, it is a genuine distributional impact rather than a handful of fluke outliers.

Other Outcomes: Though endline cases and deaths per 100,000 is the most logical outcome (see Section 2.1), we show in Appendix A.6 the results are qualitatively similar if we instead use raw counts and the log of endline cases and deaths per 100,000.²¹

Measurement Error in COVID-19 Prevalence: One inevitable challenge to any study of COVID-19 is that the true number of cases far exceeds reported cases. One strength of our design is that rainfall is unlikely to be correlated with local testing capacity, making it unlikely that our result is spuriously driven by non-classical measurement error. However, we cannot rule out that counties with larger outbreaks are more aggressive in testing. Then any variation that reduces COVID-19 cases rates, be it rainfall or a hypothetical randomized controlled trial, would find accentuated impacts. We acknowledge that this caveat applies to our study as it does to any other.

Local Policy Response: One concern is that even if rainfall is exogenous, local governments might respond to either social distancing or (more likely) rising numbers of cases by instituting their own emergency orders or lockdowns. Our estimates might reflect not just the initial shock to social distancing but the policy response triggered by that shock. Although such a response is possible, it is likely to be a countervailing response. Local officials would likely loosen restrictions wherever case counts are low and vice-versa.²² That would, if anything, bias our estimates towards zero. Nevertheless we show in Appendix A.8 that controlling for a dummy for whether the county has any policy restriction by the end of the 14 day horizon of our regressions does not change the results.

Direct Impact of Weather: Some news reports and health experts have observed that warmer countries (e.g. Singapore) have been more successful in controlling outbreaks than more temperate ones (e.g. the U.S. and Western Europe). That has led to a theory that temperature may directly affect virus transmission (e.g. Sajadi et al., 2020). If the weather directly affects transmission it could violate the single-channel assumption needed for a valid instrument.

We find no evidence for a link between transmission and temperature on the last weekend in our county-level results. Regardless, all of our specifications control for temperature, making it unlikely to be driving our results. Some reports have also suggested humidity may separately affect transmission.²³ Though the evidence for this is limited, we test for whether humidity is driving the results. If the impact of rainfall on cases and deaths were through its correlation with humidity rather than its impact on social distancing, we would expect that the reduced-form impact of rainfall on cases and deaths would vanish after controlling for humidity. But we show in Appendix A.7 that the reduced-form coefficient is essentially unchanged.²⁴ Other links are possible but not yet well substantiated. It is possible that sunlight, through ultraviolet radiation, reduces virus spread. If that is true it would bias our estimates towards rainfall increasing the number of COVID-19 cases.

That said, we cannot categorically rule out that rainfall has some unanticipated impact or interaction with the environment. Given what is currently known about the virus and the nature of our own results, we believe these effects to be second-order compared to the direct impact on human behavior.

4. Results: Value of deaths averted by early distancing

Table 2 reports our estimates of the per-capita value of deaths averted by a marginal increase in early social distancing. We multiply estimates of deaths per 100,000 from Table 1.C by 4 measures of the value of a statistical life and rescale to units of dollars per percentage point reduction in people leaving home. We report point estimates (in **bold**) together with the lower and upper

²¹ To be precise, we estimate a Poisson Maximum Likelihood estimator using Eq. 2 as the link function. Unlike simply taking the log, the Poisson estimator is consistent even though endline cases and deaths equal zero in many counties (Silva and Tenreiro, 2006).

²² Brzezinski et al. (2020) find that states where people are already social distancing of their own accord are less likely to impose a lockdown.

²³ Luo et al. (2020) is one example, though they actually find that the correlation between humidity and transmission is ambiguous.

²⁴ Since we only have humidity data for 60% of the sample, controlling for it directly in all specifications (as we do with temperature) would be too costly for precision.

Table 2
Per-capita value of marginal earlier distancing (14-day horizon) .

VSL Measure	Spec. 1			Spec. 2			Spec. 3		
1. EPA 2020 VSL	(-81.02	159.33	399.67)	(27.87	132.30	236.72)	(29.08	150.17	271.25)
2. Invariant VSL	(-89.71	176.42	442.55)	(30.86	146.49	262.12)	(32.20	166.28	300.35)
3. Constant VSLY	(-37.72	74.19	186.10)	(12.98	61.60	110.22)	(13.54	69.92	126.30)
4. Inverse-U VSLY	(-70.13	137.92	345.96)	(24.13	114.52	204.91)	(25.17	129.99	234.80)
	(Lower 95% CI			Estimate	Upper 95% CI)				

Note: We report the per capita dollar value of earlier distancing. These numbers should be interpreted as the average benefit accruing to every resident of the county. Each row is a different measure of the value of a statistical life. Each set of column reports the mean (in bold) with lower and upper 95% confidence intervals based on a specification in [Table 1.C](#).

confidence intervals. We report estimates based on 4 measures of the value of a statistical life (VSL), one used by the Environmental Protection Agency (EPA) ([United States Environmental Protection Agency, 2012](#), Table 5—9), and three measures proposed in [Robinson et al. \(2020\)](#):

1. The EPA's 2020 mean VSL estimate. This estimate is the same for all deaths regardless of age and consumption expenditure.
2. An "invariant" VSL that assigns all lives equal statistical value, based on a VSL calculated by the Department of Health and Human Services. The approach is very similar to that used by the EPA, but the numbers are slightly different.
3. A measure that adjusts for age (and years of life remaining) by assuming a constant value of statistical years (VSLY). This measure is reweighted to account for the age-specific COVID-19 death rates from February to May of 2020
4. An "inverse-U" VSLY that assigns age-specific value to statistical lives based on the average consumption expenditure of that age category, which effectively assigns higher cost to deaths of middle-aged people compared to those who are younger or older. Again, this measure is reweighted to account for COVID-19's age-specific death rates

The Constant VSLY puts the cost of a COVID-19 death far below that of the other measures because deaths are so heavily skewed towards people with fewer expected years of life. Rather than take a stand on which approach is best, we report estimates based on all three (though invariant VSLs are the most common in the literature).

Though the estimates vary based on specification and choice of VSL, most suggest a similar marginal value of earlier distancing. Take for example the estimates from Specification 2 applied to the EPA's VSL for 2020. A one percent reduction in people leaving home on the weekend before the state-wide lockdown yields an average of 132 dollars per county resident, though the confidence interval includes values as low as 28 dollars and as high as 237 dollars. Estimates based on the Constant VSLY are less than half as big, but still over 60 dollars—roughly an eight-hour workday at the federal minimum wage.

The estimates in [Table 2](#) are based on deaths averted within 14 days of the statewide lockdown. In Appendix A.9, Table 13 we show that, according to Specification 2, the value of deaths averted based on any VSL increases by roughly one-third if we switch to an 18-day horizon. The EPA's VSL, for example, implies a marginal increase in early distancing yields a per capita benefit of roughly 200 dollars. This dynamic pattern explains why earlier distancing yields such large benefits. Since each additional case spawns additional cases—often at an exponential rate—a one-time reduction in cases becomes exponentially more valuable (with the caveat that, as [Fig. 3](#) suggests, the impact does eventually taper off). Marginal early distancing is an investment that yields compounding returns.

The size of these benefits are especially surprising given that they are underestimates of the total benefit of early distancing. Our estimates adjust only for the value of deaths averted. Since our data does not contain reliable information on hospitalizations, we do not account for the value of hospitalizations averted. And although we do measure the number of confirmed cases, there is no obvious way to adjust for the value of time lost during recovery. There are also reports that some COVID-19 cases have caused long-term health problems even in relatively young people.²⁵ It is too early to know how long and how severe those symptoms will be.

But we also caution that the large average benefit of deaths averted is not distributed evenly. Aside from the obvious fact that the value of a death averted largely accrues to the person spared, [Section 3.3](#) suggests that the average is largely driven by a small number of counties where earlier distancing averted major outbreaks. Table 14 in Appendix A.9 shows that there was a 2 percent reduction in the probability a county suffered 5 or more deaths per 100,000. These counties would accrue nearly 500 dollars per person by averting these deaths. But a far larger share of counties accrued little or no benefit. These results may in part explain why earlier distancing (and mandated distancing in particular) is politically fraught. Though the benefits are high in expectation, the realized benefits may be small or zero for most counties and individuals even though the realized cost is borne by everyone.

5. Discussion: Weighing costs against benefits to inform future policy

Our results suggest that a marginal increase in earlier social distancing yields large economic benefits. These estimates are large because early distancing has persistent and growing benefits even two to three weeks later, as every case averted earlier in the outbreak prevents future cases and deaths. One interpretation is that policy makers wishing to mandate social distancing would reap

²⁵ See [Gale \(9 November 2020\)](#), for example.

surprisingly large benefits from moving more quickly. This section uses our estimates to make a very rough calculation of whether these benefits outweigh the costs.

To measure the benefits we start with studies that measure the impact of lockdowns on mobility. Since mobility is a high-frequency outcome, these studies—most of which are event studies—can precisely measure the impulse response of social distancing. They are more limited in attributing a particular case or death to a specific day of social distancing. We resolve this limitation combining the impacts on distancing measured in these studies with our estimates, which link a specific episode of distancing (the last weekend before statewide lockdown) to the full trajectory of cases and deaths. This hybrid approach yields the most complete estimate of the mortality averted by a lockdown.

We translate the topline result of each study from its original form into the impact on the percentage of people leaving home over the five days after the lockdown begins. The estimates range from -3.25 percentage points (Gupta et al., 2020c) to a low of almost zero (Cronin and Evans, 2020), while Andersen (2020) finds a topline number of -2 percentage points and Dave et al. (2020) finds impacts ranging from -1.5 to -1.9 percentage points. A reasonable median estimate would assume the lockdown decreases the number of people leaving home by 1.5 to 2 percentage points. Given that our own estimates are on average based on social distancing over a two-day weekend just prior to statewide lockdown, we can multiply our estimates of deaths averted after 14 days by the impact of the lockdown to “simulate” beginning the lockdown two days earlier.²⁶ That would imply 198 to 264 dollars of per capita benefit.

Studies estimating the cost of the lockdown are somewhat more varied in their conclusions. We focus on studies of the impact on employment because reductions in consumption spending, for example, likely stem from reduced wages. Several studies find that although voluntary distancing from people afraid of the virus had substantial impacts on employment, lockdowns had little or no impact (for example: Chetty et al., 2020; Kahn et al., 2020; Rojas et al., 2020, with the caveat that the last study actually measures the impact of lifting lockdowns). Among studies that do find an impact, we rescale the percentage reduction in employment by the employment-to-population ratio and the hours of employment lost (16 hours in total, given two days of lost work and 8-hour workdays). Several studies find that the job losses were concentrated among early-stage and non-college workers (Cheng et al., 2020; Gupta et al., 2020b), suggesting these workers are likely closer to the minimum wage than not. To be safe, assume an hourly wage of double the federal minimum. Then the cost of starting a lockdown 2 days earlier ranges from roughly 2 dollars per capita (Baek et al., 2020; Gupta et al., 2020b) to 4 dollars (Cheng et al., 2020) to 14 dollars (Coibion et al., 2020).²⁷

This back-of-the-envelope estimate of the cost of an earlier lockdown is dwarfed by the benefits. Assuming a higher hourly wage for lost work—say, 3 or 4 times the federal minimum wage—would not change the conclusion. There are two reasons why benefits heavily outweigh costs. First, the benefit of earlier distancing accumulates exponentially even after the extra two days end, while it is reasonable to assume that the cost, at least for such a marginal change in policy, is limited to the two days of lost income. Second, since deaths have consequences long after the lockdown and even after the pandemic, the amortized value of a life saved is enormous relative to the per capita cost of two days of lost labor. The main counterargument is that the extra two days of unemployment might persistently lower a worker's chances of being rehired (Gregory et al., 2020). But the initial evidence suggests that rehiring of furloughed workers has been swift (Cheng et al., 2020). Nevertheless, the true utility cost of the lockdown depends in part on whether government policy keeps employers in businesses long enough to rehire, and whether unemployment benefits arrive in time to avoid causing persistent harm to households. Neither is a trivial assumption.

Likewise, though focusing on employment is a reasonable first approximation, it does not capture impacts that might be mediated through changes in asset prices or business income.²⁸ Impacts on consumption are hard to interpret because households may reduce (or increase) consumption even if their income does not change. Nevertheless, we acknowledge that our rough estimates of the cost of lockdowns likely miss some of the monetary cost. They also discount the substantial non-monetary cost of lockdowns, such as loneliness (Hamermesh, 2020), delayed health care, and increased rates of intimate partner violence (Leslie and Wilson, 2020). On the other hand, our estimates also undercount the benefits of earlier lockdowns because they do not include the benefits of hospitalizations averted, days of labor lost to sickness, and the poorly understood long-term health impacts of COVID-19 (“long COVID”).

Despite these caveats, our rough calculations suggest that when facing a future outbreak—at least of a virus that spreads and kills at a rate equal to or greater than the novel coronavirus—the benefit of a marginally earlier lockdown may outweigh the cost by a substantial margin. Given the recent experience of the U.S., however, it is worth asking whether mandating an earlier lockdown is politically feasible. Individual responses to the crisis have been heavily influenced by partisanship (Allcott et al., 2020; Andersen, 2020; Ding et al., 2020). The decision of whether to impose a lockdown may thus be politically fraught. Conditional on doing so, however, the decision of whether to begin a few days sooner is marginal and might be relatively uncontroversial. But the results of Section 3.3 suggest there are also political challenges. A leader who mandates an earlier lockdown will often have little to show because most of her slower neighbors will have no worse an outbreak. The benefit of speed—reducing the risk of a very rare, very large outbreak—will rarely be observed.

²⁶ This is only an approximation, as the median county's lockdown actually begins 3 days after the last weekend. That might imply our estimates are actually informative about benefits accrued within 17 days from starting a lockdown two days sooner.

²⁷ We calculate these figures by taking the estimated percentage reduction in employment, rescaling by the March 2020 no. of employed persons divided by the total U.S. population, and multiplying by 16 hours and 14.5 dollars per hour. The March number of employed persons is based on the seasonally adjusted estimate from the Employment Situation Summary (Bureau of Labor Statistics, 2021). The estimate of the U.S. population is based on early estimates for April 2020 from the U.S. Census Bureau (2020).

²⁸ Gupta et al. (2020c) summarizes recent work that studies the impact of the crisis and of lockdowns on consumption.

6. Final caveats and directions for future research

Aside from the caveats about the cost-benefit exercise, there are a few implicit assumptions underpinning our main results. As noted above we cannot categorically rule out that rainfall directly affects COVID-19 transmission through some as-yet unknown mechanism. But there are also several caveats that are ripe for future research.

First, our estimates are based on data from the early days of the pandemic and can best be considered informative about a disease as unknown and deadly as the novel coronavirus was during its first wave of infections. Both the actual danger and the perceived danger have at times declined (as treatments improved and vaccines became available) and also increased (as deadlier variants arose). Future work should explore whether there are similar impacts at times of lower and higher infectiousness. Second, our estimates are relatively short-run. Since private behavior may eventually respond to actual case rates (Brzezinski et al., 2020), a long-run analysis may reveal that early- and late-distancing counties ultimately converge because lower initial case rates cause people to let down their guard.

Third, the type of social distancing induced by rainfall may differ from that induced by a government order, especially depending on whether the order is pre-announced or comes as a surprise.²⁹ Future work should test whether the type of person who stays home during bad weather is more willing to visit crowded indoor areas than the type of person who obeys a lockdown. Finally, this paper exploits variation in earlier distancing but cannot say much about subsequent behavior. In particular, our results do not speak to whether an earlier lockdown allows for an earlier lifting of the lockdown. This idea is consistent with our finding that exponential growth in case rates explains why the benefit of an earlier lockdown are so large. But future work should explore whether an earlier lockdown that also ends earlier (or even has a shorter total length) also yields fewer deaths.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jhealeco.2021.102575](https://doi.org/10.1016/j.jhealeco.2021.102575)

CRediT authorship contribution statement

Ajay Shenoy: Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing, Visualization, Supervision. **Bhavyaa Sharma:** Investigation, Data curation, Writing – review & editing. **Guanghong Xu:** Investigation, Data curation, Writing – review & editing. **Rolly Kapoor:** Investigation, Data curation, Writing – review & editing. **Haedong Aiden Rho:** Investigation, Data curation, Writing – review & editing. **Kinpritma Sangha:** Conceptualization, Writing – review & editing.

References

- Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M.D., 2020. A Multi-Risk SIR Model With Optimally Targeted Lockdown. Technical Report. National Bureau of Economic Research.
- Acemoglu, D., Johnson, S., 2006. Disease and Development: The Effect of Life Expectancy on Economic Growth. Technical Report w12269. National Bureau of Economic Research.
- Allcott, H., Boxell, L., Conway, J.C., Ferguson, B.A., Gentzkow, M., Goldman, B., 2020. What Explains Temporal and Geographic Variation in the Early US Coronavirus Pandemic? Technical Report w27965. National Bureau of Economic Research.
- Altindag, O., Erten, B., Keskin, P., 2020. Mental health costs of lockdowns: evidence from age-specific curfews in Turkey. Forthcoming at American Economic Journal: Applied Economics.
- Amuedo-Dorantes, C., Kaushal, N., Muchow, A.N., 2020. Is the Cure Worse than the Disease? County-Level Evidence from the COVID-19 Pandemic in the United States. Technical Report. National Bureau of Economic Research.
- Andersen, M., 2020. Early evidence on social distancing in response to COVID-19 in the united states.
- Anderson, T.W., Rubin, H., 1949. Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of Mathematical Statistics* 20 (1), 46–63.
- Andrews, I., Stock, J.H., Sun, L., 2019. Weak instruments in instrumental variables regression: theory and practice. *Annu Rev Econom* 11, 727–753.
- Angrist, J., Imbens, G., 1995. Identification and estimation of local average treatment effects. National Bureau of Economic Research Cambridge, Mass., USA.
- Angrist, J.D., Imbens, G.W., Rubin, D.B., 1996. Identification of causal effects using instrumental variables. *J Am Stat Assoc* 91 (434), 444–455.
- Baek, C., McCrory, P.B., Messer, T., Mui, P., 2020. Unemployment effects of Stay-at-Home orders: evidence from high frequency claims data. *Rev Econ Stat* 1–72.
- Bailey, M.J., Goodman-Bacon, A., 2015. The war on poverty's experiment in public medicine: community health centers and the mortality of older americans. *American Economics Review* 105 (3), 1067–1104.
- Baird, S., Hicks, J.H., Kremer, M., Miguel, E., 2016. Worms at work: long-run impacts of a child health investment*. *Q J Econ* 131 (4), 1637–1680. doi:[10.1093/qje/qjw022](https://doi.org/10.1093/qje/qjw022). <https://academic.oup.com/qje/article/131/4/1637/2468871>.
- Bargain, O., Aminjonov, U., 2020. Trust and compliance to public health policies in times of COVID-19. *J Public Econ* 192, 104316.
- Barrios, J.M., Benmelech, E., Hochberg, Y.V., Sapienza, P., Zingales, L., 2021. Civic capital and social distancing during the COVID-19 pandemic. *J Public Econ* 193, 104310.
- Bartscher, A.K., Nationalbank, D., Seitz, S., Sieglösch, S., Slotwinski, M., Wehrhöfer, N., 2021. Social capital and the spread of covid-19: insights from european countries. *J Health Econ* 102531.
- Bazzi, S., Fiszbein, M., Gebresilasse, M., 2021. “Rugged individualism” and collective (in) action during the COVID-19 pandemic. *J Public Econ* 195, 104357.
- Bhattacharya, J., Gathmann, C., Miller, G., 2013. The Gorbachev anti-Alcohol campaign and russia's mortality crisis. *American Economic Journal: Applied Economics* 5 (2), 232–260.
- Bleakley, H., 2003. Disease and development: evidence from the American South. *J Eur Econ Assoc* 1 (2–3), 376–386.
- Bleakley, H., 2010. Malaria eradication in the americas: a retrospective analysis of childhood exposure. *Am. Econ. J. Appl. Econ.* 2 (2).
- Brzezinski, A., Deiana, G., Kecht, V., Van Dijke, D., 2020. Covid economics. *Covid Economics* 7, 115. Preprint
- Bulfone, T.C., Malekinejad, M., Rutherford, G.W., Razani, N., 2020. Outdoor transmission of SARS-CoV-2 and other respiratory viruses: A Systematic review. *J. Infect. Dis.* 223 (4), 550–561. doi:[10.1093/infdis/jiaa742](https://doi.org/10.1093/infdis/jiaa742).

²⁹ For example, state and county lockdowns have sparked protests and political opposition, while a rainy weekend presumably would not.

- Bureau of Labor Statistics, 2021. Labor force statistics from the current population survey <https://www.bls.gov/web/empsit/cpseea01.htm>, Accessed 19 April.
- Center for Systems Science and Engineering (Johns Hopkins University), 2020. Covid-19 data repository. <https://github.com/CSSEGISandData/COVID-19>.
- Cheng, W., Carlin, P., Carroll, J., Gupta, S., Rojas, F.L., Montenegro, L., Nguyen, T.D., Schmutte, I.M., Scrivner, O., Simon, K.I., Wing, C., Weinberg, B., 2020. Back to Business and (Re)employing Workers? Labor Market Activity During State COVID-19 Reopenings. Technical Report w27419. National Bureau of Economic Research.
- Chetty, R., Friedman, J., Hendren, N., Stepner, M., Others, 2020. How did covid-19 and stabilization policies affect spending and employment? a new real-time economic tracker based on private sector data. NBER working paper (w27431).
- Chin, A.W.H., Chu, J.T.S., Perera, M.R.A., Hui, K.P.Y., Yen, H.-L., Chan, M.C.W., Peiris, M., Poon, L.L.M., 2020. Stability of SARS-cov-2 in different environmental conditions. *The Lancet Microbe* 1 (1), e10. doi:10.1016/S2666-5247(20)30003-3.
- Clay, K., Lewis, J.A., Severini, E.R., Wang, X., 2020. The Value of Health Insurance during a Crisis: Effects of Medicaid Implementation on Pandemic Influenza Mortality. Technical Report w27120. National Bureau of Economic Research.
- Cohen, J., Saran, I., 2018. The impact of packaging and messaging on adherence to malaria treatment: evidence from a randomized controlled trial in Uganda. *J Dev Econ* 134, 68–95.
- Coibion, O., Gorodnichenko, Y., Weber, M., 2020. Labor Markets During the COVID-19 Crisis: A Preliminary View. Technical Report w27017. National Bureau of Economic Research.
- Corona Data Scraper, 2020. Corona data scraper. <https://coronadatascraper.com/#timeseries.csv>.
- Couture, V., Dingel, J., Green, A., Handbury, J., Williams, K., 2020a. Device exposure indices. <https://github.com/COVIDExposureIndices/COVIDExposureIndices>.
- Couture, V., Dingel, J., Green, A., Handbury, J., Williams, K., 2020b. Location exposure indices. <https://github.com/COVIDExposureIndices/COVIDExposureIndices>.
- Courtemanche, C., Garuccio, J., Le, A., Pinkston, J., Yelowitz, A., 2020. Strong social distancing measures in the United States reduced the COVID-19 growth rate. *Health Aff* doi:10.1377/hlthaff.2020.00608.
- Cronin, C.J., Evans, W.N., 2020. Private Precaution and Public Restrictions: What Drives Social Distancing and Industry Foot Traffic in the COVID-19 Era? Technical Report w27531. National Bureau of Economic Research.
- Dave, D., Friedman, A.I., Matsuzawa, K., Sabia, J.J., 2020. When do Shelter-in-Place orders fight COVID-19 best? Policy heterogeneity across states and adoption time. *Econ Inq*.
- Dave, D., Friedman, A.I., Matsuzawa, K., Sabia, J.J., Safford, S., 2020. Were urban cowboys enough to control COVID-19? Local shelter-in-place orders and coronavirus case growth. *J Urban Econ* 103294.
- Ding, W., Levine, R., Lin, C., Xie, W., 2020. Social Distancing and Social Capital: Why U.S. Counties Respond Differently to COVID-19. Technical Report w27393. National Bureau of Economic Research.
- Endo, A., Abbott, S., Kucharski, A.J., Funk, S., et al., 2020. Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside china. *Wellcome Open Research* 5 (67), 67.
- Fairlie, R.W., 2020. The Impact of COVID-19 on Small Business Owners: The First Three Months after Social-Distancing Restrictions. Technical Report. National Bureau of Economic Research.
- Fowler, J.H., Hill, S.J., Obradovich, N., Levin, R., 2020. The effect of stay-at-home orders on COVID-19 cases and fatalities in the United States. *medRxiv* doi:10.1101/2020.04.13.20063628.
- Gale, J., 9 November 2020. Covid long haulers describe the devastating aftereffects of the disease. *Bloomberg*. <https://www.bloomberg.com/news/features/2020-11-09/coronavirus-long-haulers-tell-us-their-symptoms-and-the-aftereffects-of-disease>.
- Glaeser, E.L., Gorbach, C., Redding, S.J., 2020. JUE insight: how much does COVID-19 increase with mobility? evidence from New York and four other U.S. cities. *J Urban Econ* 103292.
- Greenstone, M., Nigam, V., 2020. Does social distancing matter? University of Chicago, Becker Friedman Institute for Economics Working Paper (2020-26).
- Gregory, V., Menzio, G., Wiczer, D.G., 2020. Pandemic Recession: L or V-Shaped? Technical Report w27105. National Bureau of Economic Research.
- Gupta, S., Montenegro, L., Nguyen, T.D., Lozano-Rojas, F., Schmutte, I.M., Simon, K.I., Weinberg, B.A., Wing, C., 2020. Effects of social distancing policy on labor market outcomes. NBER Working paper (w27280).
- Gupta, S., Montenegro, L., Nguyen, T.D., Rojas, F.L., Schmutte, I.M., Simon, K.I., Weinberg, B.A., Wing, C., 2020. Effects of Social Distancing Policy on Labor Market Outcomes. Technical Report w27280. National Bureau of Economic Research.
- Gupta, S., Simon, K., Wing, C., 2020. Mandated and voluntary social distancing during the COVID-19 epidemic: A review. Technical Report. National Bureau of Economic Research, Cambridge, MA.
- Hamermesh, D.S., 2020. Lock-downs, Loneliness and Life Satisfaction. Technical Report w27018. National Bureau of Economic Research.
- Hamory, J., Miguel, E., Walker, M.W., Kremer, M., Baird, S.J., 2020. Twenty Year Economic Impacts of Deworming. Technical Report w27611. National Bureau of Economic Research.
- Institute for Health Metrics and Evaluation, 2020. Covid-19 projections. <https://covid19.healthdata.org/united-states-of-america>. Accessed 17 April 2020.
- Jayachandran, S., Lleras-Muney, A., Smith, K.V., 2010. Modern medicine and the twentieth century decline in mortality: evidence on the impact of sulfa drugs. *American Economic Journal: Applied Economics* 2 (2), 118–146. doi:10.1257/app.2.2.118.
- Kahn, L., Lange, F., Wiczer, D., 2020. Labor demand in the time of COVID-19: evidence from vacancy postings and UI claims. NBER working paper (w27061).
- Knittel, C.R., Ozaltun, B., 2020. What Does and Does Not Correlate With COVID-19 Death Rates. Technical Report. National Bureau of Economic Research.
- Kupferschmidt, K., 2020. Why do some COVID-19 patients infect many others, whereas most don't spread the virus at all? *Science*. 19 May
- Lee, D.S., McCrary, J., Moreira, M.J., Porter, J.R., 2021. Valid t-ratio Inference for IV. Technical Report. National Bureau of Economic Research.
- Leslie, E., Wilson, R., 2020. Sheltering in place and domestic violence: evidence from calls for service during COVID-19. *J Public Econ* 189, 104241.
- Luo, W., Majumder, M.S., Liu, D., Poirier, C., Mandl, K.D., Lipsitch, M., Santillana, M., 2020. The role of absolute humidity on transmission rates of the COVID-19 outbreak. *medRxiv* doi:10.1101/2020.02.12.20022467.
- Madestam, A., Shoaib, D., Veuger, S., Yanagizawa-Drott, D., 2013. Do political protests matter? Evidence from the tea party movement*. *Q J Econ* 128 (4), 1633–1685. doi:10.1093/qje/qjt021.
- Manson, S., Schroeder, J., Van Riper, D., Ruggles, S., 2019. IPUMS National historical geographic information system: Version 14.0 [Database]. Minneapolis, MN: IPUMS doi:10.18128/D050.V14.0.
- McLaren, J., Wang, S., 2020. Effects of Reduced Workplace Presence on COVID-19 Deaths: An Instrumental-Variables Approach. Technical Report w28275. National Bureau of Economic Research.
- Menne, M.J., Durre, I., Vose, R.S., Gleason, B.E., Houston, T.G., 2012. An overview of the global historical climatology network-daily database. *J. Atmos. Oceanic Technol.* 29 (7), 897–910.
- Miller, D., Martin, M.A., Harel, N., Kustin, T., Tirosh, O., Meir, M., Sorek, N., Gefen-Halevi, S., Amit, S., Vorontsov, O., Wolf, D., Peretz, A., Shemer-Avni, Y., Roif-Kaminsky, D., Kopelman, N., Huppert, A., Koelle, K., Stern, A., 2020. Full genome viral sequences inform patterns of SARS-cov-2 spread into and within Israel. *medRxiv* doi:10.1101/2020.05.21.20104521.
- Pei, S., Kandula, S., Shaman, J., 2020. Differential effects of intervention timing on COVID-19 spread in the United States. *medRxiv* doi:10.1101/2020.05.15.20103655.
- Robinson, L.A., Sullivan, R., Shogren, J.F., et al., 2020. Do the benefits of COVID-19 policies exceed the costs? exploring uncertainties in the Age-VSL relationship. *Risk Anal.* risa.13561, doi:10.1111/risa.13561.
- Rodriguez, O.R., 2020. Fast decisions in bay area helped slow virus spread. *Associated Press*. <https://apnews.com/10c4e38a0d2241daf29a6cd69d8d7b43>, 19 April.
- Rojas, F.L., Jiang, X., Montenegro, L., Simon, K.I., Weinberg, B.A., Wing, C., 2020. Is the Cure Worse than the Problem Itself? Immediate Labor Market Effects of COVID-19 Case Rates and School Closures in the U.S. Technical Report w27127. National Bureau of Economic Research.
- SafeGraph, 2020a. Safegraph social distancing metrics, version 2. <https://docs.safegraph.com/docs/social-distancing-metrics>.
- SafeGraph, 2020b. Safegraph weekly patterns metrics, version 1. <https://docs.safegraph.com/docs/weekly-patterns>.

- Sajadi, M.M., Habibzadeh, P., Vintzileos, A., Shokouhi, S., Miralles-Wilhelm, F., Amoroso, A., 2020. Temperature and latitude analysis to predict potential spread and seasonality for COVID-19. Preprint, Available at SSRN 3550308.
- Silva, J.S., Tenreiro, S., 2006. The log of gravity. *Rev Econ Stat* 88 (4), 641–658.
- Spiegel, M., Tookes, H., 2020. Business restrictions and COVID-19 fatalities1. *Covid Economics* 20.
- Stock, J.H., Wright, J.H., 2000. GMM with weak identification. *Econometrica* 68 (5), 1055–1096.
- Thunstrom, L., Newbold, S., Finnoff, D., Ashworth, M., Shogren, J.F., 2020. The benefits and costs of flattening the curve for COVID-19. Available at SSRN 3561934.
- United States Environmental Protection Agency, 2012. Regulatory impact analysis for the final revisions to the national ambient air quality standards for particulate matter. <https://www3.epa.gov/ttn/ecas/regdata/RIAs/finalria.pdf>.
- U.S. Census Bureau, 2020. Demographic analysis uses birth and death records, international migration data and medicare records to produce a range of population estimates as of April 1, 2020 <https://www.census.gov/library/stories/2020/12/census-bureau-provides-population-estimates-for-independent-evaluation-of-upcoming-census-results.html>, 15 Dec.
- Ward, C.J., 2014. Influenza vaccination campaigns: is an ounce of prevention worth a pound of cure? *American Economic Journal: Applied Economics* 6 (1), 38–72. doi:10.1257/app.6.1.38.
- Welsch, D.M., 2020. Do masks reduce COVID-19 deaths? a county-level analysis using IV. *Covid Economics* 57 (13), 20–45.
- Ziedan, E., Simon, K.I., Wing, C., 2020. Effects of State COVID-19 Closure Policy on NON-COVID-19 Health Care Utilization. Technical Report w27621. National Bureau of Economic Research.